

Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System

S. Bhaskaran, S. D. Desai, P. J. Dumont,
B. M. Kennedy, G. W. Null, W. M. Owen Jr.,
J. E. Riedel, S. P. Synnott, R. A. Werner

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Abstract

NASA's New Millennium Program involves a series of missions whose primary purpose is to demonstrate the feasibility of new technologies for spaceflight. Deep Space 1, the first mission in the New Millennium Program, will demonstrate, among other things, an Ion Propulsion System to provide thrust, and an autonomous onboard navigation system to guide the spacecraft. The mission plan is to flyby an asteroid, Mars, and a comet using these and other new technologies.

The onboard navigation system, in order to be as self-contained as possible, uses images of asteroids taken by the spacecraft's camera as its sole data type in determining the spacecraft's trajectory. These images are clustered at intervals varying from hours to a week depending on the phase of the mission, with up to 12 different asteroids sighted per cluster. The images are then incorporated into a least-squares filter at periodic intervals to estimate spacecraft orbit parameters. The orbit determination solutions are in turn used by the navigation system to compute maneuvers required to guide the spacecraft to its targets. Since this navigation strategy has never been used in flight before, it is important to perform pre-launch assessments of its performance. This is accomplished by the use of Monte Carlo simulations which drive the navigation software with a truth model of the spacecraft trajectory and the observables. The truth model simulates realistic errors for both which are expected in flight, and individual realizations of these errors are drawn from random samplings of the errors with provided statistics. This technique is used to analyze the first leg of the mission, the flyby of the asteroid McAulliffe. The results indicate that, under nominal conditions, the combined orbit determination/maneuver computation strategy is capable of navigating the spacecraft to a safe flyby. In addition, the propulsive events required are within the abilities of the hardware to provide.

INTRODUCTION

Standard navigation techniques for interplanetary spacecraft involve the use of a combination of radio (two-way coherent Doppler and ranging) data, obtained by tracking the spacecraft using antennas at JPL's Deep Space Network (DSN) tracking stations, augmented by optical data from an onboard camera during encounters. This combination of data is is very accurate and has been used to successfully navigate spacecraft to all planets in the solar system except Pluto, and to three asteroids. However, in order to fully realize NASA's vision of the future of deep-space exploration, with multiple small, inexpensive spacecraft roaming the solar system, it is desirable to automate some or all of the processes required for interplanetary missions, including navigation. It is possible

to fully automate the navigation process by eliminating the radio data and using an onboard camera to triangulate the spacecraft's position by observing multiple solar system bodies. In this system, the data would be processed by an onboard filter to obtain the complete spacecraft ephemeris, from which maneuvers could be planned and performed to achieve the desired targeting conditions. Such a system is being developed for JPL's Deep Space 1 asteroid/comet flyby mission, the first in the New Millennium Program series of missions. The purpose of this paper is to analyze the performance of the orbit determination (OD) and maneuver targeting links of the DS-1 autonomous navigation system. Specifically, the ability of the system to deliver the spacecraft to its first target is assessed.

THE MISSION

The New Millennium Program is a recent program instituted by NASA with the primary purpose of demonstrating new technologies for future space missions. Its ambitious goal is to fly a series of missions, both Earth orbiting and interplanetary, each testing technologies which have not been proven in flight conditions and which have dramatic potential of enabling missions which could not be flown previously or of lowering the cost of space flight. The hope is that the missions will prove these technologies so that future science oriented missions can use them without incurring the cost or risk of flying a new technology. More information on the New Millennium Program can be found on its web site at <http://nmp.jpl.nasa.gov>.

Deep Space 1 (DS-1) is the first of the interplanetary missions of the New Millennium Program. In addition to autonomous navigation, other primary technologies being demonstrated include the first use of an ion propulsion system for trajectory control, an advanced solar array for power, and low-mass imaging system named MICAS (Miniature Integrated Camera and Spectrometer) (see Ref. 1 for a more detailed description of all the technologies to be validated, and Ref. 2 for an overview of all aspects of autonomous navigation). The mission itself will be launched onboard a Delta 7326 rocket between July 1 and July 31, 1998, perform a close (less than 20 km) flyby of the asteroid 3352 McAuliffe on January 20, 1999, receive a gravity assist from the planet Mars on April 2000, and then finally rendezvous with comet West-Kohoutek-Ikemura (W-K-I) in early June of 2000 at a distance of about 500 km. The main science return will come from high resolution imaging of the asteroid and comet during their respective flybys using the MICAS camera.

ION PROPULSION SYSTEM

Perhaps the most important aspect of the DS-1 mission in terms of its impact on navigation is the use of an Ion Propulsion System (IPS) engine. Unlike chemical propulsion systems which burn for short periods of time at very high thrust, the IPS produces very little thrust but is capable of burning for very long periods of time. Ionized xenon is accelerated by passing it through a charged grid before exiting out of the nozzle. The resulting thrust is on the order of milliNewtons, with specific impulses reaching values in the thousands of seconds (as compared to 200-400 seconds for chemical rockets). The thrust can be throttled by varying the voltage on the grids; for DS-1, the IPS has about 100 throttle levels, with a thrust range of 20 to 90 mN. Since the power is generated from the solar arrays, the maximum achievable thrust depends on the distance to the sun.

The characteristics of an IPS trajectory are different from those using chemical engines. Trajectories using chemical engines have long coast periods punctuated by near-instantaneous velocity changes at given times to achieve course corrections. IPS trajectories, on the other hand, are characterized by long thrusting periods of weeks to months, interspersed with coast arcs when the IPS is shut off. For DS-1, the thrusting periods have the dual purpose of providing enough energy to the spacecraft to reach its targets, and correcting launch injection, OD, and maneuver execution errors to achieve the desired targeting conditions. More details on the latter will be described below.

Designing the low-thrust reference trajectory for DS-1 is a complicated process. Briefly, the first step is to compute an optimal trajectory which takes the spacecraft from its launch injection conditions to the targets. The trajectory is optimized by finding the set of control parameters (the

right ascension α and declination δ of the thrust pointing vector, and the duration of thrusting) which achieves the targets with a minimum amount of fuel usage. Since this process is dependent on several factors, including the launch date and available power from the solar array, the nominal trajectory is constantly being revised as new data (especially about solar array performance) is received. To analyze the OD performance, we used a single reference trajectory whose characteristics should not deviate greatly from the final one flown. For the current design of the mission trajectory, the nominal IPS thrusting period begins on July 16 (15 days after launch), and ends on September 4, 1998. Prior to this period, the IPS will be used primarily for calibrating engines during its initial checkout phase. After this period, the nominal thrusting phase, or mission burn period, is over, and the IPS will only be turned on for trajectory correction maneuvers (TCMs).

If the launch injection were perfect and the IPS thrusting in exactly the designed direction and magnitude, then the mission burn would be sufficient to achieve the targets and no TCMs would be needed. In reality, of course, errors in these and other factors cause trajectory deviations, and corrections are necessary. Thus, the onboard navigation system will be used to periodically check the position and velocity of the spacecraft and correct the thrust parameters as needed. This is accomplished in the following manner. At seven day intervals during cruise, the IPS is shut down for a period of about 12–16 hours while the spacecraft slews to take sightings of up to 12 asteroids (each of these thrust/shutdown segments is referred to as a planning cycle). These observations are used to compute an OD solution to get the current spacecraft state. This state is mapped forward to the next encounter, and if the deviation from the desired encounter condition is large enough, a linearized course correction consisting of adjustments to the α and δ of the thrust vector during subsequent planning cycles, and the duration of the final mission burn segment, is computed. After the mission burn is over, the OD solution at the end of each planning cycle will be used to support TCM opportunities every few weeks. These TCMs will consist of a single IPS burn at a computed direction and duration. In the final 30 days prior to asteroid encounter, the planning cycles will have shorter durations of variable length, and the final 4 TCMs will be performed using the hydrazine based reaction control system (RCS) thrusters. These thrusters are normally used for attitude control, but due to the short time remaining before encounter, it was decided that IPS burns may require too much time to implement. Table 1 lists the times and types of maneuver opportunities for this reference trajectory. Note that both the IPS and RCS TCMs come in pairs several hours apart. This is to allow for vectorization of the maneuver, whereby if a computed thrust vector is in a direction which violates a spacecraft attitude constraint, it is broken into two segments in allowable directions whose vector sum is equal to the original. A complete description of the linear correction strategy used to correct the mission burns and compute TCMs is given in Ref. 3. Assuming that the IPS performs reasonably close to its specifications, the linear correction strategy will suffice. However, if there are very large deviations in the IPS performance from its design, or if frequent outages occur during mission burns, a redesign of the reference trajectory will be done on the ground and uplinked to the spacecraft.

ORBIT DETERMINATION

Orbit determination is the process by which the spacecraft's state (position and velocity) and other parameters relevant to the trajectory, such as nongravitational accelerations acting on the spacecraft, are estimated. In order to keep this process as self contained onboard the spacecraft as possible, the only data used to obtain an OD solution are images taken of solar system bodies (asteroids in this case) by the MICAS camera. In principle, the procedure to obtain a simple position fix of the spacecraft in heliocentric space using asteroid sightings is extraordinarily simple. A single sighting of an asteroid places the spacecraft along the line-of-sight (LOS) to that asteroid. Observing a second asteroid at the same time will deterministically fix the three-dimensional heliocentric position of the spacecraft, provided the ephemerides of the sighted asteroids and the inertial pointing direction of the camera are known. In practice, however, two simultaneous sightings are not practical with one camera, and instead, a series of LOS fixes are taken of several asteroids. For DS-1, the number of sightings taken during a given observation window of opportunity is limited by the amount of

Table 1: Maneuver Schedule for Nominal DS-1 Trajectory

Maneuver ID	Maneuver Type	Date	Time to Asteroid Encounter
0	Mission Burn	July 16, 1998 15:00:00	188 days
1	Mission Burn	July 23, 1998 15:00:00	181 days
2	Mission Burn	August 1, 1998 15:00:00	172 days
3	Mission Burn	August 8, 1998 15:00:00	165 days
4	Mission Burn	August 15, 1998 15:00:00	158 days
5	Mission Burn	August 22, 1998 15:00:00	151 days
6	Mission Burn	August 29, 1998 15:00:00	144 days
7	Mission Burn	September 5, 1998 15:00:00	137 days
8	Mission Burn	September 12, 1998 15:00:00	130 days
9	IPS TCM	September 19, 1998 15:00:00	123 days
10	IPS TCM	September 19, 1998 22:00:00	123 days
13	IPS TCM	October 10, 1998 15:00:00	102 days
14	IPS TCM	October 10, 1998 22:00:00	102 days
18	IPS TCM	November 7, 1998 15:00:00	74 days
19	IPS TCM	November 7, 1998 22:00:00	74 days
23	IPS TCM	December 5, 1998 15:00:00	46 days
24	IPS TCM	December 5, 1998 22:00:00	46 days
28	IPS TCM	December 31, 1998 20:53:46	20 days
29	IPS TCM	January 1, 1999 03:53:46	19 days, 17 hours
31	IPS TCM	January 1, 1999 03:53:46	19 days
31	IPS TCM	January 10, 1999 20:53:46	10 days
32	IPS TCM	January 11, 1999 03:53:46	9 days, 17 hours
33	IPS TCM	January 15, 1999 20:53:46	5 days
34	IPS TCM	January 16, 1999 03:53:46	4 days, 17 hours
35	IPS TCM	January 18, 1999 20:53:46	2 days
36	IPS TCM	January 19, 1999 03:53:46	1 days, 17 hours
37	RCS TCM	January 19, 1999 20:53:46	1 day
38	RCS TCM	January 19, 1999 21:13:46	1 days, 23 hours, 40 minutes
39	RCS TCM	January 20, 1999 08:53:46	12 hours
40	RCS TCM	January 20, 1999 09:13:46	11 hours, 40 minutes
41	RCS TCM	January 20, 1999 14:53:46	6 hours
42	RCS TCM	January 20, 1999 15:13:46	5 hours, 40 minutes
43	RCS TCM	January 20, 1999 17:53:46	3 hours
44	RCS TCM	January 20, 1999 18:13:46	2 hours, 40 minutes

time it takes to slew the spacecraft from one asteroid to another; an upper limit of 12 is anticipated. Several clusters of sightings are then incorporated into a least-squares filter to obtain an OD solution. The accuracy of this type of data is dependent on several factors, including the angular separation, brightness, and distance to the imaged asteroids, the resolution of the camera, the ability to pinpoint the location of the asteroid in the camera frame (centerfinding), the accuracy of the camera pointing information, and the knowledge of the asteroid ephemerides. These factors will be addressed in the following sections. For clarity, the term “beacons” are used to denote the asteroids used solely for triangulation, while “target” refers to the objects being encountered (asteroid McAulliffe and comet W-K-1 for DS-1).

The Camera System

The MICAS camera system actually has two imaging devices, one a standard charge-coupled Device (CCD), and the other an experimental active pixel sensor (APS) array. Of these, it is anticipated that the autonomous navigation (autonav) system will primarily use the CCD because of its larger field-of-view (FOV). Use of the APS by the autonav system will be limited to the final 30 minutes prior to encounter when the CCD image will be oversaturated. Both are connected to a 677 mm focal length telescope. The CCD has a 1024×1024 pixel array, giving a total FOV of 0.8° , or about 14 mrad. Each pixel therefore has an angular resolution of $13 \mu\text{rad}$.

Image Processing

The image processing link forms the core of the autonav system. Its primary purpose is to predict the locations of beacons at given times, determine the center of the asteroid in the camera frame (a process known as centerfinding), and compute the associated pointing of the camera boresight. The ability of the navigation system to perform autonomously hinges on its ability to accurately perform the centerfinding and ensuring that bad data do not corrupt the solution.

Computing predicts of beacon asteroids is the simplest of these procedures. A list of beacon asteroids to observe as a function of time for the entire mission is stored onboard the spacecraft, along with ephemerides of all the beacons (more will be said about the choice of beacons later). At predetermined times, the current spacecraft trajectory is differenced with the nominal ephemeris of given beacon to get the relative pointing vector. This information is then passed to the spacecraft attitude control system (ACS) which slews the spacecraft to the correct orientation at the correct time and shutters the picture with the provided exposure length.

Because of its importance, the centerfinding algorithms (and the associated pointing solution) used during cruise when asteroids are distant point sources have had the most testing. The details of these procedures have been documented in Refs. 4 and 5; only a brief description will be given here. The algorithms are a modification of similar ones used for the Galileo mission, both onboard the spacecraft and on the ground. They use a pattern matching technique to filter out unwanted bright spots and locate the asteroid and known stars in the camera FOV. From experimental results (see Refs. 4 and 5), the algorithms are capable of determining the location of the asteroid relative to the stars to a precision of 0.1 pixels.

For computing the pointing direction of the camera boresight, an initial guess of the values are needed. This is provided by the ACS system, which uses a wide FOV star tracker for attitude knowledge and control. The accuracy of the pointing available from ACS is about 0.3 mrad prior to alignment calibrations, and 0.1 mrad after. If at least three stars are visible in the CCD image, however, the pointing information can be improved by computing a least-squares fit to the pointing (α and δ of the boresight, and the twist around the boresight) using the ACS values as an initial guess. Assuming 0.1 pixel centerfinding ability, the pointing can be determined to a few μrad . If fewer than three stars are available, then the accuracy is degraded. Analysis has shown that three or more stars will be available during cruise. Encounter navigation requires new data types because the extended target body is very bright (usually about magnitude 2-3 per pixel) and because very near encounter the target image fills the camera field of view. When stars are in the field but the

contrast between the stars and target body exceeds the camera dynamic range, then “flash-mode” observations are made by alternating short exposures of the target body and longer exposures to bring up stars; camera pointing is determined from the star exposures and interpolated to the time of the target body exposure. When the range to the target is sufficiently small (an hour from closest approach for McAulliffe), then “starless” observations of the target are processed using the camera pointing values obtained from the star tracker.

The star catalog used by autonav contains 221,594 stars that lie within 30° of the ecliptic and have a catalogued visual magnitude of 10.50 or brighter. The positional data for the stars are taken from the highly accurate catalogs produced by the European Space Agency’s Hipparcos satellite.

For purposes of evaluating the OD, an observation uncertainty, σ_o , of 0.1 pixel was used for the beacon observations, which represents the current best estimate of the centerfinding accuracy for distant, unresolved asteroids. As the spacecraft nears encounter, however, the target asteroid becomes resolved and the pattern matching centerfinding algorithm cannot be used. Instead, a simple brightness centroiding on the asteroid is done. Because the asteroid has an unknown shape, this method can only determine the brightness center, and the true center is unknown. The error is potentially as large as the radius of the asteroid, so, the data are deweighted accordingly. The uncertainty used is the angular extent of the body in the camera FOV, converted to pixels:

$$\sigma_o = \frac{\tan^{-1}(R/\rho)}{13 \times 10^{-6}}, \quad (1)$$

where

$$\begin{aligned} R &= \text{the assumed radius of the asteroid,} \\ \rho &= \text{the range to the asteroid.} \end{aligned}$$

Asteroid Ephemerides

An implicit assumption in the use of triangulation asteroids for orbit determination is that the heliocentric positions of the asteroids at the time of the observation is known exactly. In fact, this is not really the case; knowledge of the orbits of the 5,000 or so numbered asteroids are known to different accuracies. The larger and/or brighter objects which have been tracked for longer periods of time have orbital accuracies in the tens of km, while the smaller and dimmer objects which have not been observed as much are known to within only several hundreds of km.

To properly account for the ephemeris errors, the orbits of the asteroids used for triangulation would have to be estimated along with the spacecraft trajectory in the OD filter. However, this would greatly increase the complexity of the filter since there are over 80 beacons. To keep the onboard OD algorithm simple, therefore, asteroid ephemeris errors are ignored. Instead, by using up to 12 asteroids per observation set, we rely on simple averaging to remove these errors during cruise.

Encounter presents a special case. For a given camera and centerfinding ability, the accuracy of an observation is directly proportional to the distance to the asteroid. During encounter, the target is several orders of magnitude closer than the beacons so the power of its observations overwhelms the information provided by the beacons. The result that the spacecraft’s target relative state is accurate to the level of the data, but its heliocentric state estimate is skewed by an amount roughly equivalent to the ephemeris error present in the target’s orbit. This is an acceptable consequence though, since it is the target relative, not the heliocentric, state which is important for targeting and visually tracking the object during the flyby.

In order to minimize the adverse effects of ignoring the asteroid ephemerides, a ground campaign is underway to improve the orbits of some asteroids. About 80 asteroids have been identified as probable beacons for the current DS-1 trajectory; 31 of these are being observed from JPL’s Table Mountain Observatory with the expectation that their orbits can be improved by a factor of 3 or 4. Of particular importance are the flyby targets, McAulliffe and W-K-I. The current prediction is

that, assuming the observations are successful, McLaughlin's orbit uncertainty can be improved from its current value of 127 km, 50 km, and 60 km in the radial, transverse, and normal directions, respectively, by about a factor of 3.

Beacon Asteroid Selection

One non-autonomous portion of the navigation function is the selection of the beacons used for triangulation. This procedure, referred to as the picture planner, is done on the ground and the results stored onboard before launch. The picture planner propagates the spacecraft state and asteroid states using either conic elements or numerical integration. For each planned weekly triangulation session, it searches for acceptable observing opportunities by examining observation characteristics for the lowest-numbered 5000 asteroids and selecting the subset of asteroids which produce the best combined accuracy in the local instantaneous spacecraft state determination. These computations take into account camera sensitivity, full well, system noise, and dynamic range. Observation geometry conditions constraining beacon selection include beacon brightness, beacon distance, solar phase angle, spacecraft pointing constraints, camera measurement accuracy, star background (at least two suitably bright stars are required), and star-relative smear of the beacon during the computed exposure time (the cross-correlation can tolerate only 1-2 pixels of star-relative smear). Closer asteroids provide better observation accuracy provided that the star-relative smear is acceptably small. Attitude control performance parameters such as absolute pointing accuracy (about 0.1 of the CCD field) and expected limit cycle "kick velocity" (about 3 pixels/sec) are also used in the picture planning computations. Camera exposure time and pointing can be adjusted to provide the best astrometric measurement accuracy for each observation opportunity. For each selected asteroid the output includes observation epoch, asteroid identification, exposure time, and the few-hour effective span for which the prediction is valid. The trajectory file for the beacon asteroids will typically contain 100-200 asteroids. For encounter, the picture planner output is referenced to the encounter time and the onboard navigator then updates the absolute observation times using its latest encounter time determination.

Dynamical Equations and Filtering

In general, the process of determining orbital state parameters of an interplanetary spacecraft is a nonlinear one. However, the process can be considerably simplified by linearizing the problem, which amounts to solving for deviations of the orbit parameters about a reference trajectory rather than the orbit parameters themselves. This allows powerful methods of linear estimation theory to be applied, resulting in more stable solutions. This does require, though, that initial guesses to the state parameters be available to generate the reference orbit.

The second-order equations of motion used to generate a reference trajectory can be written as two first-order equations:

$$\dot{\mathbf{r}} = \mathbf{v} \quad (2)$$

$$\dot{\mathbf{v}} = -\frac{\mu_s}{r^3}\mathbf{r} + \sum_{i=1}^{n_p} \mu_i \left[\frac{\mathbf{r}_{ri}}{r_{ri}^3} - \frac{\mathbf{r}_{pi}}{r_{pi}^3} \right] + \frac{AG}{mr^3}\mathbf{r} + \frac{k}{m}\mathbf{T} + \mathbf{a}, \quad (3)$$

where

- \mathbf{r} = the heliocentric cartesian position vector of the spacecraft,
- \mathbf{v} = the heliocentric cartesian velocity vector of the spacecraft,
- \mathbf{r}_{pi} = the heliocentric cartesian position of the i th perturbing planetary body
- \mathbf{r}_{ri} = the position of the spacecraft relative to the i th perturbing body, i.e., $\mathbf{r}_{ri} = \mathbf{r}_{pi} - \mathbf{r}$
- μ_s = the gravitational constant, GM , of the sun,
- μ_i = the gravitational constant of the i th perturbing planet,

- n_p = the number of perturbing planets,
- A = the cross-sectional area of the spacecraft,
- G = the solar flux constant,
- \mathbf{T} = the thrust vector from the IPS, in Newtons,
- k = the thrust scale factor, with values between 0 and 1,
- m = the spacecraft mass, and
- \mathbf{a} = miscellaneous accelerations acting on the spacecraft.

In Eq. 3, the first term on the right hand side represents the central body gravitational acceleration from the sun, the second term is the sum of the third body gravitational acceleration from the planets (all except Pluto are used), the third is the solar radiation pressure, the fourth is the acceleration due to thrusting from the IPS, and the final term accounts for any additional unmodeled accelerations acting on the spacecraft.

The two gravitational acceleration contributions are straightforward, but the non-gravitational forces acting on the spacecraft deserve some discussion. With regard to solar radiation pressure, it is obvious from Eq. 3 that a simple spherical model for the spacecraft was used. In reality, the spacecraft's cross-sectional area is dominated by the two solar array panels, with the spacecraft bus contributing a much smaller proportion. During flight, the panels will almost always be pointed at the sun, with the bus rotating to provide thrust vector control, camera pointing, etc. Since the dominant effect is from the panels which remain more or less fixed relative to the sun, it was decided that the complexity of using a more accurate model was not needed.

Thrusting events on the spacecraft come from two sources: the IPS for mission burns and TCMs, and the RCS for attitude control and late TCMs. IPS events are explicitly accounted for in the filter via the fourth term in Eq. 3, but are handled differently depending on whether the integration is performed from a past time to the present for OD purposes (the data arc), or for predicting the state of the spacecraft at some future time (predicts). For the former, the actual thrust achieved by the IPS is not measured directly (such as with an accelerometer), but is instead indirectly computed based on measured voltages across the ion acceleration grid. At preset intervals varying from seconds to minutes, the voltage is read out for computing the magnitude of the thrust, and the spacecraft attitude at the corresponding time is also obtained from the ACS to get the thrust direction. This information is passed to the navigation system which accumulates the high rate data and, when a certain threshold in either the thrust magnitude or change in direction is reached, prints a record to a history file containing an averaged thrust magnitude and direction over that time span. This averaging minimizes the storage required to maintain history information over a long data arc. Since thrust is not directly measured, the value of thrust computed will have some uncertainty associated with it. The characteristics of the measurement error is somewhat uncertain at this time, but is expected to be within $\pm 1.5\%$ of the true value. The scale factor k in the fourth term of Eq. 3 is used to account for this measurement error and will be an estimated parameter in the filter.

RCS thrusters are used primarily for attitude control, but will also be used for TCMs near encounter. Once again, the way they are handled in the integration depends on whether the integration is over a past time or for predicts. For history information, onboard ACS software sends out thruster activity reports in terms of the velocity change, or $\Delta \mathbf{V}$, accumulated over a time span, with the minimum time span presently set to 1 second. The navigation software receives these high rate messages and compresses the data by waiting until a minimum $\Delta \mathbf{V}$ threshold is reached, after which a record of the total $\Delta \mathbf{V}$ vector at that time is written to the same history file which stores the IPS activity. Additional records are also written if a time threshold is passed, so that small $\Delta \mathbf{V}$ s which do not reach the threshold will be properly time tagged. Finally, prior to obtaining an OD solution, all remaining $\Delta \mathbf{V}$ s which have not reached either the magnitude or time threshold are written to the history file. RCS activity for attitude control and TCMs are handled in this manner without any distinction being made between the two. Unlike the IPS, though, each individual RCS event is not modeled explicitly in the filter. Instead, the fifth term in Eq. 3., the general acceleration term, is used to estimate the averaged acceleration errors over a given span caused by mismodeling of the

ΔV s caused by RCS firings.

At the time when an OD solution is needed, the integration over the data arc proceeds as follows. The history file is sorted so that the IPS and RCS propulsive activity records are time ordered. Starting from the beginning of the data arc, the integrator will proceed through the time span, stopping and restarting at each thruster event. For IPS events, the acceleration contribution of the thruster is interpolated by computing the thrust magnitude and direction over the time span in which it is active, dividing the thrust magnitude by the current spacecraft mass to get acceleration. For the RCS, the instantaneous ΔV contribution at a given time is added to the current velocity vector, and the new state is propagated forward. Although this method of stopping and starting the integrator at thrust discontinuities is time consuming, the accuracy gained is substantial and necessary to prevent filter divergence.

For predicts, the thrust value during IPS mission burns will be computed as a function of the available power from the solar arrays, which in turn is a function of the distance to the sun. During IPS TCMs, the thrust is nominally zero but will be adjusted by the maneuver software for retargeting. The adjusted value and associated duration are written to a maneuver file; this file is read by the integration routine to obtain the appropriate thrust information. The scale factor k for IPS predicts will always take a value of 1. RCS TCMs also are nominally zero, adjusted during retargeting, and written on the maneuver file. As with the history integration, the integrator for predicts will stop at RCS events on the maneuver file, add the instantaneous computed ΔV , and restart the integration. The accuracy of this method remains only if the RCS ΔV s are small (on the order of m/s) and therefore take only a short time for the thrusters to achieve; large RCS ΔV s would incur an integration error penalty. Current analysis indicates that RCS TCMs should indeed be fairly small, so this is not currently a cause for concern. Finally, attitude control events in the future are not predictable and presumably average to zero over the course of the mission. For this reason, they are not modeled and the general acceleration term in Eq. 3 is ignored for predicts.

Filter

Once the reference trajectory for the data arc is generated, the solution of the state parameters, which are corrections to the nominal values used to generate the reference, can be obtained using the techniques of epoch state batch filtering from linear estimation theory [Ref. 6]. If we define the parameters of the trajectory to be adjusted, $q(t)$, as

$$q(t) = [X(t) \ Y(t) \ Z(t) \ \dot{X}(t) \ \dot{Y}(t) \ \dot{Z}(t) \ k_1 \dots k_n \ a_x \ a_y \ a_z]^T, \quad (4)$$

where

- X, Y, Z = the cartesian position components,
- $\dot{X}, \dot{Y}, \dot{Z}$ = the cartesian velocity components,
- $k_1 \dots k_n$ = thrust scale factors, with a different scale factor estimated for each planning cycle in a data arc, and
- a_x, a_y, a_z = the components of the general acceleration vector,

then the updated trajectory, $q'(t)$, is

$$q'(t) = q^*(t) + \Delta q(t), \quad (5)$$

where the asterisk denotes the nominal trajectory, and $\Delta q(t)$ is the vector of estimated corrections (henceforth, the Δ will be eliminated in the notation for the correction vector q). If the nominal values are reasonably close to the truth, then the corrections should be linear over the batch time span, and the corrections at the epoch state, $q(t_0)$, can be linearly mapped to any other time t using the state transition matrix, Φ , as

$$q(t) = \Phi(t)q(t_0), \quad (6)$$

where,

$$\Phi(t) = \frac{\partial \mathbf{q}^*(t)}{\partial \mathbf{q}^*(t_0)}. \quad (7)$$

To get the state transition matrix values at a given time, t , note that

$$\dot{\Phi} = \frac{\partial \dot{\mathbf{q}}^*(t)}{\partial \mathbf{q}^*(t)} \frac{\partial \mathbf{q}^*(t)}{\partial \mathbf{q}^*(t_0)} = \mathbf{A} \Phi(t). \quad (8)$$

This matrix differential equation represents a set of $(9 + N_{pc})^2$ scalar first-order equations, where N_{pc} is the number of planning cycles in a data arc. The initial condition is $\Phi(t_0) = \mathbf{I}$, the identity matrix. The partial derivative matrix, \mathbf{A} is computed analytically. Many of the elements of \mathbf{A} are zero, so that only $7N_{pc} + 63$ equations are needed. These equations are integrated along with the nominal trajectory to get Φ as a function of time.

To set up the equations for the epoch state batch filter, the partial derivatives of the observations with respect to the state are needed. The observables in this case are the pixel p and line l coordinates of the beacon or target asteroid centers obtained using the centerfinding techniques described earlier. Thus, at the time of the observation, the partials matrix \mathbf{H} is

$$\mathbf{H} = \begin{bmatrix} \partial p / \partial X & \partial p / \partial Y & \partial p / \partial Z & 0_{2 \times (6 + N_{pc})} \\ \partial l / \partial X & \partial l / \partial Y & \partial l / \partial Z & 0_{2 \times (6 + N_{pc})} \end{bmatrix} \quad (9)$$

The observed (p, l) depends only on the spacecraft's position relative to the beacons at the instant the image is taken; hence the partials with respect to the velocity and acceleration components are zero. The non-zero values of \mathbf{H} can be computed analytically, and the equations for these partials are given in Ref. 7. To map these partials back to the epoch, the state transition matrix is used:

$$\tilde{\mathbf{H}} = \mathbf{H} \Phi, \quad (10)$$

where $\tilde{\mathbf{H}}$ is the observation partial matrix at epoch. Given the *a priori* covariance matrix, \mathbf{P}_0 , the observation weighting matrix, \mathbf{W} (a diagonal matrix whose elements are $1/\sigma_o^2$, with σ_o being the observation uncertainties from Eq. 1), and a residual vector, \mathbf{Y} , which are the differences between of the observed centroid values and the predicted ones computed from the nominal spacecraft trajectory, the original epoch state batch filter equations for the solution vector $\hat{\mathbf{q}}$ and the formal covariance \mathbf{P} are:

$$\hat{\mathbf{q}} = [\mathbf{P}_0 + \tilde{\mathbf{H}}^T \mathbf{W} \tilde{\mathbf{H}}]^{-1} \tilde{\mathbf{H}}^T \mathbf{W} \mathbf{Y} \quad (11)$$

and

$$\mathbf{P} = [\mathbf{P}_0 + \tilde{\mathbf{H}}^T \mathbf{W} \tilde{\mathbf{H}}]^{-1}. \quad (12)$$

In practice, however, the equivalent *U-D* factorized method is used [Ref. 8]. In this method, which was adopted to minimize round-off error and ensure stability, \mathbf{P} is expressed as the product $\mathbf{U} \mathbf{D} \mathbf{U}^T$, where \mathbf{U} is upper triangular with ones on the diagonal and \mathbf{D} is diagonal.

After an initial testing phase, the OD solution strategy to be adopted is as follows. After the first 28 days of cruise during which autonav is enabled, an OD solution is performed. Nominally, this means that four planning cycles are incorporated with 12 observations in each planning cycle, resulting in 48 observations. The *a priori* covariance matrix, \mathbf{P}_0 , for the solution is set such that the position and velocity components are effectively unconstrained, with values of 10^8 km and 100 m/s used for the 1σ uncertainties in position and velocity, respectively. The nominal scenario calls for thrusting during this period, so four thrust scale factors, corresponding to each of the 7-day planning cycles, are also estimated, with *a priori* uncertainties for each set to 5%. Finally, the *a priori* sigmas on the components of the general acceleration term are set to 3×10^{-9} km/s. These values allow the filter to freely adjust the spacecraft's initial position and velocity while constraining the thrust and accelerations to be within reasonable bounds.

Following this initial solution, solutions are performed at 7-day intervals during cruise by dropping the data from the earliest planning cycle and adding the data from the planning cycle just completed.




Figure 1: 1σ Uncertainties in Position and Velocity vs Data Arc Length

Thus, the OD is performed over a sliding window of a constant 28-day length. The number of planning cycles in this window will vary during encounter when planning cycle lengths shorten to less than 7 days, but the total number of days is always kept constant (the amount of data in each of these 28-day batches will also vary during encounter). The same values for \mathbf{P}_0 are used every time a solution is done, so effectively, each batch solution has no “memory” of a previous solution, and data information older than 28 days is lost. The nominal starting trajectory to which corrections are made, however, is the latest in that the starting values for position and velocity at a given batch epoch are the mapped values from the previous OD solution.

The rationale for using this solution strategy can be seen from the plot in Fig. 1. On the figure, the mean position and velocity formal sigmas are plotted as a function of batch length. It can be seen that the uncertainties make noticeable improvements when data from 14, 21, and 28 day batches are used, but they quickly level off afterwards. This is due to the relatively large non-gravitational accelerations acting on the spacecraft, primarily from the IPS thrusters. The noise in these accelerations hinders the mapping of information from one time to the next so that after 28 days or so, the data add little information to solve for the epoch state. For this reason, 28 days was chosen to be the optimal batch length, providing enough information to obtain a reasonable solution but not cluttering the filter with useless data. However, by using the sliding batch window approach and updating the nominal trajectory at each OD solution, the nominal trajectory is implicitly computed with information older than 28 days, after the first OD solution.

The formal uncertainty plots in Fig. 1, show that, for a typical cruise data arc, the filter can determine the spacecraft position to about 130 km in position and 0.7 m/s in velocity. It is also instructive to see how well the filter can estimate the thrust scale factors. Fig. 2 plots the formal uncertainties in the estimate of four thrust scale factors in a typical 28 day arc. In this run, the *a priori* uncertainty on the scale factor was set to 5%. It is clear from Fig. 2 that the first scale factor is poorly determined, with no improvement from the *a priori*, while the fourth is best determined (to about 2%— a little less than half of the *a priori*). In general, the later scale factor estimates will be better than the earlier, although for this particular thrust profile, the second scale factor is better determined than the third due to the orientation of the thrust vectors. The reason for the first scale factor being so poorly determined is twofold: first, the first set of observation data is taken 7 days after the epoch, and second, the epoch state is unconstrained. This results in all errors being absorbed by the state, with nothing attributed to the thrust. Conversely, the fourth planning cycle has data at both ends thereby tightly constraining the position, with the result that remaining




Figure 2: 1σ Uncertainties in Thrust Scale Factor Estimates

errors have to be absorbed by the scale factors. Even then, the improvement is not dramatic, so some care must be taken in interpreting the values of the scale factors estimated by the filter.

The complete set of dynamics and filter will be used to obtain OD solutions throughout the cruise and up to 30 minutes prior to the nominal encounter time. After this, it is expected that the processing time required with the onboard computer resources is not sufficient to permit rapid turnaround of the OD result to update pointing predicts during the flyby. For this reason, the target observations taken after Encounter (E) – 30 minutes will be brightness centroided and passed to a fast, compact 3-state filter (named the Reduced State Encounter Navigation filter, or RSEN). A version of the RSEN filter has already been developed for a similar flyby of a comet for the STARDUST mission, and a description of the algorithm and its performance is given in Ref. 9. The observations are used to update the target relative position only; the target relative velocity has been well determined at this point. The initial state for RSEN is provided at the E–30 minute point from the main navigation module. RSEN is then used primarily to maintain visual lock on the asteroid during the period surrounding closest approach; it will not be used to support further TCMs.

MONTE CARLO SIMULATION AND RESULTS

If the dynamic equations used in the filter accurately modeled the true forces acting on the spacecraft and the errors in the observations were also correctly represented, then the formal covariance obtained after filtering would accurately represent the statistics of the estimated values. This is clearly not the case however, as we have deliberately simplified the nongravitational acceleration terms and ignored some of the errors which affect the data. For this reason, Monte Carlo simulations are needed to assess the true filter performance and compare it with the formal uncertainties. For the simulations, a “truth” model of the trajectory and observations are generated and provided to the filter. For a given run, the truth model represents a single realization from a random sampling of the errors which affect that model. One hundred runs are performed, and the results evaluated by computing statistics on the difference between the known truth and the estimated values computed by the filter. The details of this process will now be described.

Trajectory Model

The trajectory model used for the truth integration is the same as in Eqn 3, with a couple of additions. These modifications are used to simulate errors in the true thrust output by the engines, and to

simulate errors in the measurement of the thrust provided to the autonav system. These simulated differences are modeled as sinusoid functions of time. If T and T' represent the commanded and true thrust magnitudes, respectively, then the magnitude for the thrust term in Eq. 3 used for the truth integration is

$$T' = T + A_T \cos[2\pi(t - t_0)/\tau_T], \quad (13)$$

where

- A_T = the amplitude of the additional thrust magnitude,
- t, t_0 = the current and epoch time, and
- τ_T = the time constant of the magnitude variation.

The right ascension α and declination δ of the truth thrust vector are similarly modeled.

The measured thrust magnitudes and directions are taken as a variations on the true values, *i.e.*,

$$T^* = T' + B_T \cos[2\pi(t - t_0)/\tau_T] \quad (14)$$

and similarly for α^* and δ^* . These measured quantities are passed to the OD filter. The measured thrust vector is broken into time ordered segments and sent to the autonav routines to be placed into the history file. Thus, the information used to integrate the trajectory in the filter will be different from the truth integration used to generate observables. This best mimics what will happen onboard the real spacecraft where the true thrust produced by the IPS will not be known to the filter.

For a given set of Monte Carlo runs, the amplitude terms in Eqs. 14 and 14 are random samples with zero mean and given standard deviation. The time constants, however, will be kept constant for a particular set of runs. The amplitude of the thrust magnitude variations is expressed as a percentage of the nominal thrust value, and the direction amplitudes are in degrees.

In addition to thrust, the initial state (position and velocity) is varied. Each run of the Monte Carlo simulation will be started with a random sample of zero mean and assumed standard deviation around the nominal initial state. In general, this is the largest error source for which the filter must solve.

Observation Model

During flight, the observables will be taken from centroiding on images of asteroids. Although the capability exists to generate simulated images to centroid, the time it takes to generate a single image precludes their use in a 100-sample Monte Carlo run. Thus, the observable generation was simplified to taking samples of the expected statistics of the observations. The process used is as follows. The true spacecraft-to-beacon vector is computed using the truth spacecraft trajectory and truth asteroid ephemerides. This vector is converted into camera coordinates, and random noise is added, with the noise having zero mean and a given standard deviation. The resulting pixel and line values are passed to the filter as the observations.

As mentioned earlier, the ephemerides of the asteroids are not perfectly known, and the error is not accounted for in the filter. This effect is simulated in the Monte Carlo runs by using a different ephemeris for the truth observable generation as compared to the nominal ephemerides used by the filter to get the computed observables. To get a precise representation of this error would require that the covariance of the ephemerides of each beacon asteroid be sampled, and this value added to the nominal ephemerides for the truth. This process is time consuming, however, so a simpler solution was used. A single number representing a crude mean of the ephemeris errors of all the beacons is used, and a random sample for the three-dimensional position error of each beacon is drawn using this value as the standard deviation and added to the nominal to get the truth. For the flyby target asteroid, though, a separate value for the uncertainty in the radial, transverse, and normal components of the ephemeris error is sampled to get the truth. Thus, the target asteroid, being a special case for evaluation, has a more realistic representation of its ephemeris error.

Evaluation of Results

The evaluation of the filter performance is done by differencing the known truth values with values obtained by the filter, and then collapsing the results for the 100 samples by computing statistics on the differences. The values used for evaluation depend on the mission phase. During cruise, maneuvers are computed using the epoch state value estimates mapped to the current time, which represents the best knowledge of the trajectory from which to plan course corrections. Thus, the cruise performance is evaluated by comparing the mapped heliocentric cartesian state from the filter with the concurrent true state. During encounter, however, the increasing power of the target asteroid data will cause the heliocentric trajectory to be adjusted to fit the target-relative data. If no target ephemeris errors were present, then the heliocentric and target-relative path would be the same. Since the simulation (and reality) will have this error, the estimated trajectory is adjusted by the filter so that it is correct relative to the target, but not necessarily in heliocentric space. For evaluation of the encounter results then, we use the true spacecraft state relative to the true target state, differenced with the estimated spacecraft position state to the nominal target state. In addition, the target-relative states are transformed into the so-called “B-plane” encounter coordinates. The B-plane is an imaginary plane centered on the flyby target and perpendicular to the incoming trajectory asymptote. It is defined by three mutually orthogonal unit vectors: \mathbf{S} , parallel to the relative incoming trajectory asymptote and normal to the B-plane; \mathbf{T} , in the B-plane and parallel to the program reference plane (Earth Mean Equator of J2000.0); and $\mathbf{R} \equiv \mathbf{S} \times \mathbf{T}$, also in the B-plane. The intersection of the incoming asymptote with the B-plane defines the vector \mathbf{B} , whose components are denoted as $\mathbf{B} \cdot \mathbf{R}$ and $\mathbf{B} \cdot \mathbf{T}$. Finally, distances in the \mathbf{S} direction are usually converted into equivalent times of flight by dividing by the hyperbolic excess velocity.

Another criterion used for evaluation is the additional $\Delta \mathbf{V}$ needed to achieve the target beyond the nominal thrusting. Recall that the nominal thrusting includes only the mission IPS burns early in the cruise; IPS and RCS TCMs are nominally zero. The combination of launch injection errors and OD and maneuver execution errors during the course of the mission cause deviations from the nominal trajectory which need to be corrected by the TCMs. For the IPS, corrections in the direction do not require additional fuel, but corrections to the duration do. Thus, the amount of change in the IPS durations required to correct the errors is a measure of performance. Similarly, statistics on the required $\Delta \mathbf{V}$ for the RCS burns are also computed and presented.

Results

The results for the nominal case assume the current best estimates for baseline error values which affect the trajectory and the observations. The following uncertainty values were used (all values are 1σ):

- Initial launch + 16 day injection errors of 5000 km in position, 0.5 m/s in velocity.
- IPS thrust magnitude execution errors of 2% of the nominal.
- IPS thrust direction execution errors of 1.0° in α and δ .
- Time constant for execution errors in magnitude and direction of ∞ (in other words, the error is a bias across the mission duration).
- IPS thrust magnitude measurement errors of 1.5% of the nominal.
- IPS thrust direction measurement errors of 0.05° .
- Time constant for measurement errors of ∞ .
- Data noise of 0.1 pixel.
- Beacon asteroid ephemeris errors of 100 km.

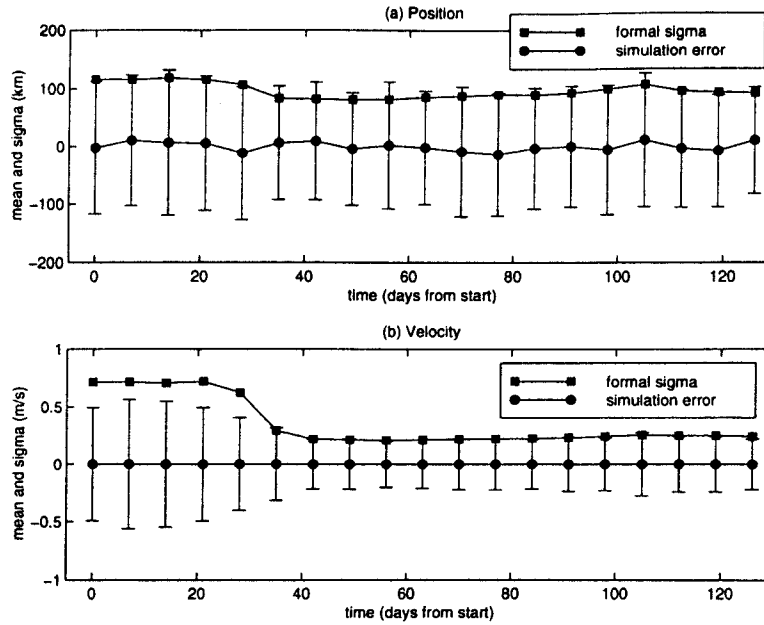


Figure 3: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Cruise

- Target asteroid ephemeris errors of 40 km, 16 km, and 20 km in the radial, transverse, and normal directions, respectively.

The cruise results are shown in Fig. 3. Plotted are the formal covariance sigmas obtained by the filter as well as the mean and standard deviation of the actual errors from the 100 Monte Carlo simulation runs. Overall, the statistics of the actual errors matched the predicted uncertainty from the filter. In position, the errors in the early and late parts of the cruise came fairly close to the formal sigma, while in the middle, the standard deviation was roughly 1.5σ . In velocity, the standard deviations never exceeded the formal sigmas and the time history of simulation statistics almost exactly matched that of the formal sigmas. In addition, since the mean of the errors was near zero, the implication is that ignoring asteroid ephemeris errors did not introduce significant biases into the estimates, and that these errors were sufficiently averaged out. The effect of the nongravitational accelerations is shown by the fact that the estimates did not improve markedly over the course of the mission. The initial position determination was good to about 120 km, and this improved to only about 95 km. Slightly more improvement was seen in the velocity error, which decreased from about 0.5 m/s to 0.2 m/s.

Although these results for the heliocentric spacecraft trajectory are not as accurate as those achievable by standard Doppler and range tracking, the advantage of using optical data becomes obvious when examining its capability of delivering the spacecraft to its target. Fig. 4 shows a plot of the mean and standard deviation of the truth minus estimated errors in the target B-plane coordinates, along with the expected 1σ uncertainty from the filter, for the final 2 days before encounter. In this case, the mean values show a bias of about 0.5 km and 0.2 km in $\mathbf{B} \cdot \mathbf{R}$ and $\mathbf{B} \cdot \mathbf{T}$, and about 1.6 second in TOF. This is caused by the systematic error of the center-of-brightness to center-of-mass offset in the observations of the extended body. Because the object is expected to be small, however, this bias is not a critical factor in choosing the flyby aimpoint. The standard deviation of the errors about the mean are similar in magnitude in the crosstrack components, and about 1.6σ in TOF.

This plot clearly indicates the ability of the optical data to determine the crosstrack target-relative position of the spacecraft. Both the expected and actual errors shrink rapidly from several km at E-2 days to sub-kilometer levels at E-3 hours. In the TOF, or downtrack, component

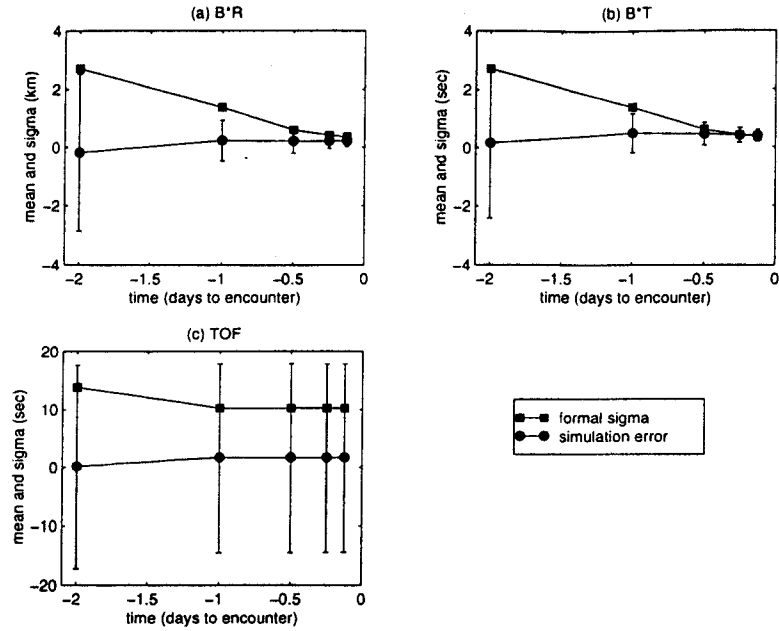


Figure 4: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Encounter

however, there is little improvement after E-2 days. For this reason, it was decided that the final four RCS TCMs only control the crosstrack errors in the B-plane, and accept the TOF control provided by the last IPS TCM. As will be described shortly, this results in considerably smaller maneuvers required by the RCS at almost no cost in delivery accuracy.

The spacecraft delivery to its flyby aimpoint is shown in Figures 5 and 6. Fig. 6 plots the target B-plane, overlain with the expected size of the asteroid, the flyby aimpoint, and the 3σ ellipsoid defining the expected uncertainty from the filter of the delivery. The scatter of dots shows the true flyby location after the E-3 hour targeting maneuver from the Monte Carlo simulation runs. Even with the half km bias in the OD results, it can be seen that the predicted subkilometer level control of the flyby aimpoint was met in about 85% of the cases. The rms of the errors was 0.8 km, and the maximum was 1.7 km. In no case was there a danger of impacting the asteroid.

The errors in the downtrack, or TOF direction, is shown as a histogram in Fig. 6. The two vertical dashed lines in this plot show the 3σ formal sigma in the TOF axis, and the histogram plots the number of samples out of the 100 which fell into a particular time bin. Once again, the majority of cases are within the formal error bounds, with only a few cases exceeding it. The maximum values are 43 seconds on the late side and 38 seconds on the early side. Overall, larger error values and sigmas are seen in the TOF axis as opposed to those in the B-plane itself due to the lack of direct information about this axis from the optical data. The errors in the TOF can be reduced only very near to encounter, when the LOS direction to the asteroid rotates perpendicular to the downtrack direction.

Clearly, the flyby results for the nominal case in all three axes are acceptable in terms of safe delivery to the target. For the primary science goal of imaging the asteroid during closest approach, however, improvements are needed. In particular, the TOF uncertainty would preclude keeping sight of the asteroid with a 0.6° FOV camera during the flyby. Thus, the RSEN filter described earlier will be used to update the pointing information. The uncertainties in the OD after the last targeting maneuver will be reduced by RSEN during the terminal approach.

The amount of change in the IPS and RCS thrust profiles needed to achieve the target conditions in the presence of nominal errors is given in Tables 2 and 3. In Table 2, the sum of all the duration changes for each sample run was tallied, and the statistics on the 100 sums were computed. The minimum duration change is negative because, in 7 samples, the final mission burn had to be




Figure 5: True B-plane Delivery Locations from Monte Carlo Simulations




Figure 6: True Time-of-flight Errors from Monte Carlo Simulations

Table 2: IPS Duration Change Statistics

Minimum Duration Change	-14.0 hours
Maximum Duration Change	57.0 hours
Mean Duration Change	22.9 ± 14.4 hours

Table 3: RCS ΔV Statistics

Minimum ΔV	0 m/s
Maximum ΔV	0.25 m/s
Mean ΔV	0.09 ± 0.04 m/s

shortened from its nominal 6 day duration, and the sum of the remaining IPS TCM durations did not exceed this decrement. At first glance, this would appear to be a benefit since less fuel is expended to reach the first target, leaving more ΔV capability for the remainder of the mission. However, since the thrust profile is optimized for the entire mission assuming a certain spacecraft mass, the heavier spacecraft may not be able to reach its second target using the nominal profile, which may prompt a redesign of the trajectory.

Table 3 shows similar statistics on the ΔV magnitude sums using the RCS engines. Here, the minimum is zero because in 1 sample, the targeting using IPS was accurate enough such that the RCS was never used for maneuvering. The worst case is only 0.25 m/s; this is easily achievable by the RCS thrusters, which have the capability of providing close to 2 m/s of ΔV .

Fig. 7 plots the ΔV statistics for each TCM. For comparison purposes, the IPS durations were converted to ΔV by applying an approximate scale factor of 10 m/s per day of IPS thrusting (in other words, an IPS duration of one day results in a ΔV of 10 m/s). The mean ΔV and its standard deviation for the 100 samples is plotted. As expected, the largest value occurred at the first IPS TCM, which made an average correction of 5 ± 2.6 m/s. In general, the earlier TCMs make larger corrections, and they are used more often. In this case for example, maneuver 9, the first TCM, was required in 92 samples, as compared to the E-1 day TCM being used in 61 samples, the E-12 hour TCM in 29 samples, the E-6 hour TCM in 19 samples, and the E-3 hour in only 5 samples.

The results from the nominal case validate the maneuver strategy of not controlling the TOF using the RCS thrusters. As a comparison, a set of Monte Carlo runs were made where all three components were targeted in the final four TCMs. These results showed an order of magnitude increase in the ΔV magnitudes, with the mean value jumping from less than 0.1 m/s to over 1.6 m/s. The maximum value in several instances hit a software limit of 2.5 m/s. The delivery in the B-plane was almost identical, and the TOF miss went from an rms value of 18.3 seconds down to 16.1 seconds. This marginal improvement in the TOF control obviously does not justify the increased fuel expenditure needed to achieve it.

CONCLUSIONS

The simulations described in this paper are the first results of performance testing on the DS-1 autonomous navigation system. This test validates the basic concept of using onboard optical sightings as the sole data type, and proves that, under certain assumptions, the system is capable of navigating a spacecraft safely to a close flyby of an asteroid. In addition, a statistical look at the additional ΔV required from the IPS and RCS engines under these assumptions was accomplished, and this revealed that the values obtained are within the capabilities of the current hardware. Finally, the simulations also served the purpose of functional testing of the components of the navigation system.




Figure 7: ΔV Statistics per Maneuver

The testing is far from over, however, and many more simulations need to be run before full confidence in the system can be established. The performance in the presence of variations in the error sources, including worst case scenarios, needs to be analyzed. In a similar vein, the software needs to be stressed to its limit to find out when and under what conditions it fails. Since an autonomous flight system needs to be exceptionally robust, these failure modes need to be identified and handled gracefully to avoid loss of the spacecraft. In addition to preparing the software, the ground testing will also prepare the analysts to handle problems and contingencies during the flight of a revolutionary method of navigation.

Acknowledgment

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

1. M. D. Rayman, D. H. Lehman, "NASA's First New Millennium Deep Space Technology Validation Flight", IAA Paper IAA-L-0502, Proceedings of the Second IAA International Conference on Low-Cost Planetary Missions, Laurel, MD, April 1996.
2. J. E. Riedel, S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen Jr., R. A. Werner, B. G. Williams, "Navigation for the New Millennium: Autonomous Navigation for Deep Space 1", Proceedings of the 12th International Symposium on Space Flight Dynamics, Darmstadt, Germany, June 1997.
3. S. D. Desai, S. Bhaskaran, W. E. Bollman, C. A. Halsell, J. E. Riedel, S. P. Synnott, "The DS-1 Autonomous Navigation System: Autonomous Control of Low-Thrust Propulsions Systems", AIAA Paper 97-38819, AIAA Guidance, Navigation and Control Conference, New Orleans, LA, August 1997.

4. R. M. Vaughan, J. E. Riedel, R. P. Davis, W. M. Owen Jr., S. P. Synnott, "Optical Navigation for the Galileo Gaspra Encounter", AIAA Paper 92-4522, AIAA/AAS Astrodynamics Conference, Hilton Head, S.C., August 1992.
5. S. Bhaskaran, J. E. Riedel, S. P. Synnott, "Autonomous Optical Navigation for Interplanetary Missions", *Space Sciencecraft Control and Tracking in the New Millennium*, E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE, pp. 32-43 (1996).
6. P. B. Liebelt, *An Introduction to Optimal Estimation*, Addison-Wesley, 1967.
7. W. M. Owen Jr., R. M. Vaughan, "Optical Navigation Program Mathematical Models", JPL Internal Document JPL-EM 314-513, August 9, 1991.
8. G. J. Bierman, "Measurement Updating Using the U-D Factorization", Proc. 1975 Con. on Decision and Control, Houston, TX, pp. 337-446.
9. S. Bhaskaran, J. E. Riedel, S. P. Synnott, "Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example", Paper AAS 97-628, AAS/AIAA Astrodynamics Conference, Sun Valley, ID, August 1997.